

1. Opening remarks and introductions

The Chairman, Don Tolmie of Los Alamos National Laboratory, opened this HIPPI-6400 meeting and thanked himself and Los Alamos for hosting this meeting. This group is constituted as both the HIPPI special working group (SWG) under X3T11, and the HIPPI Networking Forum (HNF) - Technical Committee (TC).

Don then lead a round of introductions. The list of attendees is at the end of these minutes.

2. Review / modify the draft agenda

The draft agendas were available on the web prior to the meeting. Hard copies were available at the meeting. James Hoffman of Los Alamos volunteered to take the meeting minutes. Michael McGowen asked to make a presentation on Message header format at the beginning of the meeting which was added to the agenda. The presentation and resulting discussion appears in section "4.4 Message Structure" of these minutes.

3. Review minutes of previous meetings

3.1 August 5-6, 1996, Honolulu, HI

The Honolulu minutes were reviewed and approved with the addition of text describing Craig Davidson's proposed connector pin-out and Bob Newhall's action item to review the pin-out.

3.2 Review action items from Honolulu meeting

1. Greg Chesson to provide Don Tolmie with the parallel CRC equations to be added as an annex. (Done)
2. James Hoffman to verify CRC error protection results. (Carryover)
3. Greg Chesson and SGI to specify the order that bits are fed into the LCRC calculation. (Done)
4. Greg Chesson and SGI to provide text on what actions occur as a result of a Shutdown. (Carryover)
5. Greg Chesson to provide ARP text for inclusion in HIPPI-6400-SC. (Carryover)
6. Greg Chesson to draft initial text for bridging. (Carryover)
7. Michael McGowen of Essential Communications to detail the use of the "Transtype" parameter. (Done)
8. Michael McGowen to describe the ULA and EtherType fields and their usage. (Done)
9. Hansel Collins to check the ± 200 ppm bit rate tolerance value. (Done)
10. Hansel Collins to check the $50\% \pm 5\%$ CLOCK symmetry value. (Done)
11. Hansel Collins to check the 600 mV minimum output voltage value. (Done)
12. Hansel Collins to check the 10% deterministic jitter value. (Done)
13. Hansel Collins to check the 10% random jitter value. (Done)
14. Hansel Collins to check the ≤ 0.05 ns peak jitter value. (Done)
15. Hansel Collins to check the 0.6 ns (20-80%) rise and fall time values. (Carryover)
16. Hansel Collins to determine the receiver's common mode range. (Done)
17. Craig Davidson of Raytheon E-Systems to propose the voltages and currents necessary at the connector. (Done)
18. Greg Chesson to review the second paragraph of 8.2: ACK indications on Source side. (Done)
19. Greg Chesson to propose a solution for unspecified TYPE values and whether they should be handled as errors. (Done)
20. Greg Chesson to review the Scheduled Transfer with SGI and to present Don Tolmie with the new field ordering. (Done)
21. Michael McGowen to decide whether FP can be translated. (Done)
22. James Hoffman to revise Annex C with the suggested changes. (Done)
23. Roger Ronald agreed to incorporate changes suggested at the meeting to the Admin micropacket draft. (Done)
24. Bob Newhall and SGI to review the proposed connector pin-out from Craig Davidson. (Carryover)

25. Don Tolmie to update HIPPI-6400-PH Rev 0.4 with the changes agreed to at the Honolulu meeting. (Done)
26. Don Tolmie to find a sizable whacking stick to be applied liberally to anyone failing to complete their action items. (Done?)

4. Review HIPPI-6400-PH changes since last meeting (reference HIPPI-6400-PH Rev 0.5)

4.1 Minor changes first and then return to unresolved issues in 4.8

Changed "Message" to be "one or more micropackets" – accepted.

Added an acronym for ULP – accepted.

Added text into figure 3 clarifying MSB orientation – accepted.

Clarified ECRC coverage as "up to this point" in many areas of the document – accepted.

Revised 4.10 and removed open issues – accepted.

Changed 6.2 VC Message length restrictions to account for two header micropackets – accepted. A request was made by Michael McGowen to allow unscheduled large datagram traffic on VC3. The group decided the issue needs more investigation as this will greatly impact network congestion on VC3. Michael McGowen took an action item to start an email discussion on this topic.

Revised 6.3.2 and required Null micropackets to be inserted when idle – accepted with the removal of text in the last sentence.

In 6.4 added "most recent" to speed RSEQ replies – accepted. The group also wanted some mention of dealing with stomped micropacket's RSEQ in either 6.4 or 6.6.2.1.

Clarified note 1 in 11.2 describing disparity range – accepted.

In 13.3 combined two alike bullets into bullet three – accepted.

4.2 LCRC specification and parallel equations (pages 12-13, 43-46)

Added input bit-ordering for LCRC in 6.6.2 – accepted with request for reference to table 5.

In 6.6.2.1 and 6.6.2.2 changed wording to "shall" requiring specific stomp code implementation and deleted the second note. Also add that a stomped micropacket will not "consume" a TSEQ value – accepted.

Added annexes C.3 and C.4 describing the parallel equations for generating the CRC's – to be reviewed by the group and any errors reported.

In 6.6.3, added requirement to log ECRC_Source_Error on the first occurrence of a bad ECRC in a Message.

4.3 Error processing (pages 26-28, 37)

In 9.1, the credit timeout value was deemed an extreme lack of forward progress mechanism and changed to 2 s.

In 9.2, logging an RSEQ_Missing_Error was added in the second paragraph, and logging a RSEQ_Out_Of_Range_Error was added in the last paragraph – accepted.

In 9.4, the Retry_Error was renamed Retransmission_Error. Text was added about programming the number of errors before shutdown – accepted. Henry Brandt noted that electrical interference may require a much larger programmable successive retransmission counter.

In 10.1.2, the last two error cases were deleted and the single remaining case was re-worded – accepted and presented the action item that Dave Parry or Jim Davis should review the text.

The Unspecified TYPE error = x'0'-x'7' was moved to 10.1.4 as it applies at the link layer. The error log must also record the undefined TYPE and the micropacket will be treated as a Null micropacket.

In 10.2.1, the statement that no errors are logged was removed. The group decided to treat all unspecified TYPE's in the range of x'8' - x'F' as TYPE Data. This will allow encapsulated headers to be placed after the first micropacket with TYPE Header, but not break legacy switches. All unspecified TYPE values will be logged. A single error counter will be used for any unspecified TYPE values encountered.

In 10.2.2, the possible VC numbers were changed from 0-3 to 1-2 since Admin micropackets only travel on VC1 and VC2 – accepted.

In 10.2.5, the stall timeout value should only be caused by a broken or extremely slow originating source (or possibly multiple large-granularity micropacket interleaving switches). The timeout value was changed to 1 ms. It was suggested that Dave Parry review the changed stall and credit timeouts.

A note about “1/4” was pulled out of the 14.1 text and added to table 8. Table 8 was updated to match the errors listed throughout the text – accepted. Timeout values were also changed in the table as noted above.

4.4 Message Structure (pages 14-17, 42)

Renumbered 6.7 into separate clause 7 for Messages – accepted.

Added text to clause 7 requiring that only one Message may be in-progress for each VC – accepted.

Padding a micropacket with zeros was added as discussed in Dallas. John Renwick objected to requiring padding at the ULP as an implementation specific requirement.

Combined MAC and Scheduled header tables into figure 10 – accepted.

Michael McGowen proposed a new HIPPI-6400 MAC Header format closely mirroring RFC 1042 (shown below) . Addressing issues and needs appear in section 5.2 of this document. The new Header impacts the protocol in several ways:

- Unrestricted addressing requires switches to route based off an 48-bit address only (the logical address field, the registry, and extra ULA bits disappear.)
- A restricted addressing method must be developed for switches already designed to use logical addresses.
- Two Ethertypes must be acquired to carry HIPPI-FP and “Scheduled Transfer”. (and raw?)
- Easier translation to Ethernet especially when carrying IP (at least easier than before). Bridges can use well known algorithms.

D_ULA (msb)			
D_ULA		S_ULA (msb)	
S_ULA			
M_LEN			
802.2 LLC			ORG
DSAP	SSAP	CTL	
ORG	ORG	EtherType	
Opcode	Op_Flags	M_Count	
D_Port		S_Port	

All further issues concerning the new MAC Header that weren’t discussed at the meeting will be left for the HIPPI Reflector or the next meeting.

Added normative Annex B defining how the ULA’s are coded. Since normative annexes must come before informative annexes, the informative annexes were renumbered accordingly. Due to the new MAC Header format, Annex B will be removed in future revisions.

4.5 Scheduled Transfers (pages 17-25, 47-51)

Figure 11 was changed to show the Schedule Header in the second micropacket of all Scheduled Transfer messages – accepted.

The individual operation semantics in 8.2 and 8.3 were changed. Unused fields were removed and unchanging parameters were called out once. The intent is to increase readability and avoid duplication in the text – accepted.

Tables 3 and 4 were extensively changed to match the operation semantics in 8.2 and 8.3 (removed unchanging parameters.) – accepted with changes noted below.

James Hoffman presented Scheduled Transfer open issues on two transparencies, so the issues could be argued without worrying about actual text or formatting. The decisions are presented below:

- The 16-bit Opcode was split into 6 bits of Opcode and 10 bits of Operation Flags.
- The M_flags and S_flags fields were removed as their functionality has been moved to other areas.
- The VC assignment communicated during the RTS_Response was moved into the RTS operation.
- The group decided to keep the R_limit field which was renamed to Op_limit, placed in M_count, and a zero value declares don’t care. Devices exceeding the Op_limit may have extra

Operations dropped by the Final Destination. Note there is no flow control method on port setup operations. An Op_limit value will be communicated in both the RQP and RQP_Response to maintain the full duplex nature of the operations.

- Zero was declared the well known port for Virtual Connection operations. Greg Chesson took an action item to investigate port mapping for HIPPI-6400 Virtual Connection operations.
- Having a Source_Offset communicated in the RTS operation was removed.
- A Destination may have as many in-progress Scheduled Transfers as it can handle. This requires some rewording of the document for clarity.
- It was challenged whether Blocksize should be an integral of the Final Destination Buffer size or the maximum Message size. Upon review using the maximum Message size allows better Block transmission granularity.
- The group was presented with the decision at the Hawaii meeting to require integral power-of-two buffer sizes and confirmed the selection.
- The T_len field will carry the Transfer length in the RTS_Response as indicated in the RTS operation.

The new Operation Flags field will carry the following flags:

- First - First Message of a Block
- Concatenate - combine Bufx and Offset as a 64-bit buffer index
- Notify - notify ULP on receipt of entire Block
- Persistent - retain memory for future Transfers
- Reject - Reject RQP or RTS operation (sent in RQP_Response or RTS_Response)
- Update - This operation updates parameters established in a previous similar operation (such as R_limit)
- 2 VC bits for the RTS operation

Jim Pinkerton requested additional changes from those listed above:

- The L, Last, bit was removed as it is not needed to signify the last Block of a Transfer (as Blocks may be out of order), or signal the end of a Block (already have M_count = 0).
- Text should be added to the standard requiring Messages within a Block to arrive in order and note that Blocks may arrive out of order.
- The T_id bit should be split into 8 bits of Destination ID and 8 bits of Source ID and exchanged and used in the same manner as a port. (The port field is not good enough as it will nominally be associated with the upper

processing layer, not the device involved.) The new fields provide a faster decoding capability.

- A 32-byte ULP payload will now accompany all RTS operations. A 64-bit ULP payload will end the second micropacket of RQP, CTS, RTS_Response, CTS/RTS_Response, and RQP_Response.

The EtherType originally placed in the S_type now goes into the B_num during RQP operations.

The DATA_ACK operation signals a higher layer protocol acknowledging the Block number listed in B_num and all previous Blocks. The DATA_ACK should be initiated by a higher layer protocol acknowledging complete reception and acceptance of the listed Blocks.

It was requested that the Abort operation be renamed "End" since it ends the Transfer whether in-progress or an unlimited Transfer.

It was noted that Scheduled Transfers would not be able to cross routers but should go across bridges.

4.6 Addition of CLOCK_2 signal (pages 32-33, 36)

Added the free-running CLOCK_2 signal to Figures 13 and 15. Added text defining CLOCK_2 to 15.1 and 15.2 – accepted with minor wording changes.

Silicon Graphics noted that their CLOCK_2 signal will not be phase aligned to other signals but will have constant phase.

The group noted that a 500 MHz clock may be harder to detect than other signals, but also realized there is no other method of implementing discrete deskew logic without a constant running clock.

4.7 Copper interface specifications (pages 38-40)

The group reviewed the new copper interface sections and Hansel Collins provided missing values in section 16.2 presented below.

Maximum high-level output voltage = 2.69 V
 Minimum high-level output voltage = 2.20 V
 Maximum low-level output voltage = 0.28 V
 Minimum low-level output voltage = 0.0 V
 Maximum rise time = 300 ps
 Minimum rise time = 100 ps
 Maximum fall time = 300 ps
 Minimum fall time = 100 ps

The cable specifications in 16.5 currently require values based on a per meter basis, but the group decided to specify on a per cable basis. A “within pair skew” value of 500 ps was added to 16.5.

Hansel Collins presented a filter network that he and Ron Nikel had worked on for passive equalization of 50 meter cables. The group concluded that further study was needed on the filter especially when considering much shorter cables.

Hansel Collins also presented a skew and jitter budget calculated using worst case linear addition.

<u>Data Pattern</u>	<u>Present</u>	<u>Future</u>
Dependent Jitter	440 ps	180 ps
Delay Line Res.	198 ps	160 ps
Delay Line Jitter	0 ps	12 ps
Bit Time Distortion-Driver	100 ps	80 ps
Bit Time Distortion-Receiver	0 ps	30 ps
Transmit Clock Jitter	0 ps	50 ps
Receiver Induced Data Jitter	0 ps	30 ps
Receiver Induced Clock Jitter	0 ps	30 ps
Coupled Noise Data Jitter	0 ps	0 ps
Coupled Noise Clock Jitter	0 ps	0 ps
SuMAC Jitter and Skew	100 ps	100ps
Total Jitter and Skew	838 ps	672 ps

The group noted that someone needs to review European grounding specifications, possibly with Roger Cummings. Grounding will be a big concern with many vendors and users due to strict regulations at many industry sites.

4.8 Resume discussions on items deferred in 4.1

The previous discussions seemed to hit most of the document and no new logical issues were pursued due to time and energy constraints.

5. HIPPI-6400-SC

5.1 Admin micropackets

The group was concerned about McGowen’s new Header format that no longer uses ULA’s. It was decided that the current Admin micropacket scheme would allow the logically switched E-Systems switch to assign ULA’s based on the “restricted addressing” methods. ULA assignment would be based off the Admin micropacket draft.

Roger Ronald’s Admin micropacket draft was passed out at the meeting. The operations were reviewed and refined. Roger Ronald continued his action item

to update the Admin micropacket draft with suggested revisions.

6. HIPPI-6400 MIB

Von Welch agreed to update the counters in the MIB to conform with the latest decisions.

7. Other “Open Issues” that haven’t been covered

7.1 Teleconference with the optical group (2pm Thursday)

A portion of the optical group started a teleconference on 2pm on Thursday and reported to the HIPPI-6400 group around 4pm. Stan Swirhun agreed to provide minutes of the teleconference. The optical groups will make presentations at the October X3T11 meeting and the HIPPI-6400 group will select a ferule. A matching connector will be selected in December.

The meeting covered electrical specifications, optical specifications, and connector and mechanical specifications.

The optical group asked for the same time slot they occupied at the Hawaii meeting for the October meeting, (from 8 am to 9 am and 3 pm to 7 pm on Tuesday, October 8.)

Hansel Collins commented on the divergent low level output voltages used to couple to optical devices in the market. The optical group was asked for preferences and Stan Swirhun noted both PECL 3.5 V and PECL LVDS were prevalent.

8. Low-cost optical interconnect (3 pm - 7 pm Thursday)

Some new attendees arrived so introductions were done again, and thanks again extended to Los Alamos for hosting the meeting.

8.1 Presentations

Jonathan Thatcher of IBM presented the current Jitney designs and applicability towards HIPPI-6400. It was noted that the current Jitney uses 10 fibers per connector and that a large volume dealer using 12 fibers would be needed to bring down costs on a 12 fiber unit.

8.2 Work planning

The group decided to forego any immediate work on Jitney based HIPPI-6400. Prices and connector hardware is currently out of reach of the HIPPI-6400 market, but may be considered once Jitney has been in production a while.

9. Future meeting schedule

9.1 October 7-8, 1996, St. Petersburg Beach, FL

During the X3T11 October plenary week, the following HIPPI meetings are scheduled:

Monday, October 7 -

9 AM - 9 PM — HIPPI-6400

Tuesday, October 8 -

9 AM - 10 AM — HNF Plenary

10 AM - 3 PM — HIPPI-TC General and -6400

3 PM - 7 PM — HIPPI-6400 Optical (pick a ferrule)

7 PM - 9 PM — HIPPI-6400 Copper (tentative)

The location is the Tradewinds Hotel in St. Petersburg Beach, FL. Charles Brill and AMP are the host. (See the meeting announcement on the web page at <http://www.cic-5.lanl.gov/~det/> for further details.)

9.2 November 6-7, Phoenix, AZ

Wednesday, November 6 -

1 PM - 9 PM — HIPPI-6400

Thursday, November 7 -

9 AM - 3 PM — HIPPI-6400

3 PM - 9 PM — Copper/Fiber?

The location is the Courtyard by the Marriot Hotel near the Phoenix Sky Harbor Airport. Chris Olson and Lockheed Martin/Loral are the host. (See the meeting announcement on the web page at <http://www.cic-5.lanl.gov/~det/> for further details.)

9.3 Future meetings

The schedule for the rest of 1996 is listed below. The Plenary meetings include HIPPI-6400, an HNF plenary, and all other HIPPI items. The Interim meetings cover just HIPPI-6400 items.

1996 -

Dec 2-3 Plenary Minneapolis, MN IBM

The following 1997 X3T11 plenary week dates are firm, but some locations and hosts are still fluid.

1997 -

Jan 8-9	Interim	Phoenix, AZ	Lockheed
Feb 3-4	Plenary	San Jose, CA	Sun
Mar 5-6	Interim	??	Berg
Apr 7-8	Plenary	Palm Springs, CA	Brocade
Jun 9-10	Plenary	Seattle, WA	Boeing
Aug 4-5	Plenary	Honolulu, HI	Hitachi
Oct 6-7	Plenary	Tucson, AZ	FSI
Dec 1-2	Plenary	Orlando, FL	DPT

10. Review action items

All of the following action items apply to HIPPI-6400.

1. James Hoffman to verify CRC error protection results.
2. Greg Chesson and SGI to provide text on what actions occur as a result of a Shutdown.
3. Greg Chesson to provide ARP text for inclusion in HIPPI-6400-SC.
4. Greg Chesson to draft initial text for bridging.
5. Hansel Collins to check the 0.6 ns (20-80%) rise and fall time values.
6. Michael McGowen to begin VC Message size restrictions discussion via email.
7. Greg Chesson to review counter size of SuMAC Retransmission_Error counter and the need for both contiguous retransmission and total retransmission error counters.
8. Dave Parry to review the changed values for stall timeout and credit timeout.
9. Dave Parry or Jim Davis to review the new text of 10.1.2: Check Received TSEQ.
10. Michael McGowen to look at bridging and address self discovery concerns with the new MAC Header and start a discussion on email for any unresolved issues.
11. Greg Chesson to review port mapping for HIPPI-6400 Scheduled Transfers.
12. Greg Chesson to look into IP over HIPPI-6400, focusing on the Scheduled Header.
13. Roger Ronald to update the Admin micropacket draft.
14. Don Tolmie to update HIPPI-6400-PH Rev 0.5 with the changes agreed to at the Albuquerque meeting.

11. Adjournment

The group was released around 5 pm after many attendants had left to fly home.

12. Attendance: (at main HIPPI-6400 meeting)

Michael Griffin	3M
John Renwick	Ascend
Jeff Young	Cray Research
Bob Willard	Digital Equipment Corp
Michael McGowen	Essential Communications
Jim Mott	Essential Communications
Bob Pearson	Essential Communications
Randy Hardy	Harris
Francois Gaullier	Hewlett-Packard
Henry Brandt	IBM
Chris Olson	Lockheed Martin
Mike Boorman	Los Alamos National Lab
Gene Dornhoff	Los Alamos National Lab
Andy DuBois	Los Alamos National Lab
Dave DuBois	Los Alamos National Lab
James Hoffman	Los Alamos National Lab
Ian Philp	Los Alamos National Lab
Don Tolmie	Los Alamos National Lab
Fred Templar	NASA Ames/Sterling
Von Welch	NCSA
Joe Parker	Optivision
Craig Davidson	Raytheon E-Systems
Roger Ronald	Raytheon E-Systems
Greg Chesson	Silicon Graphics
Hansel Collins	Silicon Graphics
Jim Pinkerton	Silicon Graphics

13. Attendance: (at HIPPI-6400 Low-cost optical meeting)

Michael Griffin	3M
John Renwick	Ascend
Bob Willard	Digital Equipment Corp
Jim Mott	Essential Communications
Bob Pearson	Essential Communications
Randy Hardy	Harris
Henry Brandt	IBM
Kenneth Jackson	IBM
Dan Kuchta	IBM
Jonathan Thatcher	IBM
Chris Olson	Lockheed Martin
Gene Dornhoff	Los Alamos National Lab
James Hoffman	Los Alamos National Lab
Ian Philp	Los Alamos National Lab
Don Tolmie	Los Alamos National Lab
Von Welch	NCSA
Joe Parker	Optivision
Craig Davidson	Raytheon E-Systems
Roger Ronald	Raytheon E-Systems
Greg Chesson	Silicon Graphics
Hansel Collins	Silicon Graphics